

GENESYS[®]

This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

Genesys Mobile Services Deployment Guide

Multi-site Deployment

5/5/2025

Multi-site Deployment

For a multi-site deployment, you must consider the following questions:

- How are API requests from mobile devices going to be routed?
- How is the mobile/web client going to retry the request in case of network failure?
- How will the client find the addresses for the global load balancer or site load balancer?
- How is the telephony network going to route the call to the number the client is receiving through the data API call?
- Is the call origination direction user-originated or user-terminated?

The general recommendation for user originated scenarios is to deploy an independent GMS cluster per site, and control call routing by ensuring that no intersecting pools of incoming phone numbers are configured. This way, if the client is requesting an access number through the data API call, it will be routed to Site A by the data network (internet), and then the call initiated by this client will land on the telephony switch/gateway located on (or associated with) Site A.

The matching API call will first be performed against the GMS cluster located on Site A. If Site A fails, then Site B is tried. This provides a good level of site isolation and simplifies maintenance, as well as day-to-day operations. Contact centers will be able to distribute telephony traffic between sites by controlling data API traffic (splitting between sites or directing to one site only) using standard IP load balancers, DNS records, and so on.

Using an approach of separate GMS clusters helps to avoid complex problems of "split brain" scenarios, where typically, some type of manual intervention is required either after the split, or later on when connection between sites is restored and clusters must be reconciled back together.

For user terminated scenarios, where call matching is not required, deploying a single GMS cluster across sites may provide some benefits. Still, it is not clear whether those benefits will outweigh the risks and be more beneficial as compared to the simplicity of a dedicated site cluster approach.

The following diagram shows major components of the solution and connectivity between them through local and wide area network segments. Possible network and/or component failures are numbered and described in more detail in the table below.



Failure Condition (see diagram)	Resource Locking Scenario	Two Separate GMS Clusters: One on Site A and One on Site B		
Create Service/ Get Access Number	Match ixn Scenario			
	Same Site	Cross Site		
1 or 2	Global number, no locking.	Client has to retry the other site.	Not affected.	Request fails on the local site and has to be retried on the other site.
	Two global numbers – one per site, no locking.	Client has to retry the other site.	Not affected.	Will never happen because the site specific number was given out.
	Pool of numbers with locking.	Client has to retry the other site.	Not affected.	Request fails on the local site and

Failure Condition (see diagram)	Resource Locking Scenario	Two Separate GMS Clusters: One on Site A and One on Site B		
				has to be retried on the other site.
3	Global number, no locking.	Not affected.	Not affected.	Request fails. Retrying the other site is not possible because the GMS cluster cannot be reached on another site.
	Two global numbers – one per site, no locking.	Not affected.	Not affected.	Will never happen because the site specific number was given out.
	Pool of numbers with locking.	Not affected.	Not affected.	Request fails. Retrying the other site is not possible because the GMS cluster cannot be reached on another site.
4	Global number, no locking.	Not affected.	Request fails. Retrying the other site is not possible because separate clusters are being run.	
	Two global numbers – one per site, no locking.	Not affected.		
	Pool of numbers with locking.	Not affected.		
5	Global number, no locking.	Not affected.	Request does not fail. Another GMS instance will handle the calls.	
	Two global numbers – one per site, no locking.	Not affected.		
	Pool of numbers with locking.	Not affected.		
6	Global number, no locking.	Not affected.	Request does not fail. Another ORS instance will handle the calls.	
	Two global numbers – one per site, no locking.	Not affected.		
	Pool of numbers with locking.	Not affected.		
7	Global number, no locking.	Not affected.	Request fails if there is no URS backup set in the configuration. Retrying the other site is not possible because separate clusters are being run.	
	Two global numbers – one per site, no locking.	Not affected.		
	Pool of numbers	Not affected.		

Failure Condition (see diagram)	Resource Locking Scenario	Two Separate GMS Clusters: One on Site A and One on Site B		
	with locking.			
8	Global number, no locking.	Not affected.	Not affected.	Request fails. Retrying the other site is not possible because the GMS cluster cannot be reached on another site.
	Two global numbers – one per site, no locking.	Not affected.	Not affected.	Will never happen because the site specific number was given out.
	Pool of numbers with locking.	Not affected.	Not affected.	Request fails. Retrying the other site is not possible because the GMS cluster cannot be reached on another site.